



Multi-FPGA HPC Applications - MuFPGA -

Michael Laß

Tobias Kenter, Heinrich Riebler, Christian Plessl

High-Performance Computing Group
Paderborn Center for Parallel Computing

Project Group Presentation, 18th July, 2022



Paderborn
Center for
Parallel
Computing

High-Performance Computing Group

High-Performance Computing Group

- Research in HPC methods and architectures
- Focus on accelerator technologies / heterogeneous computing
- Unique in its activity in the field of FPGA accelerated HPC
- Close connection to the Paderborn Center for Parallel Computing (PC2)



Paderborn
Center for
Parallel
Computing

HPC Applications: Modeling and Simulation



climate prediction



mechanical structure simulation



protein folding

HPC Applications: Optimization



airflow optimization



fuel combustion optimization



evolved antenna
(designed with
evolutionary algorithm)

HPC Applications: Analytics



DNA sequence analysis (bio informatics)



social networks



astrophysics (sun spot activity)

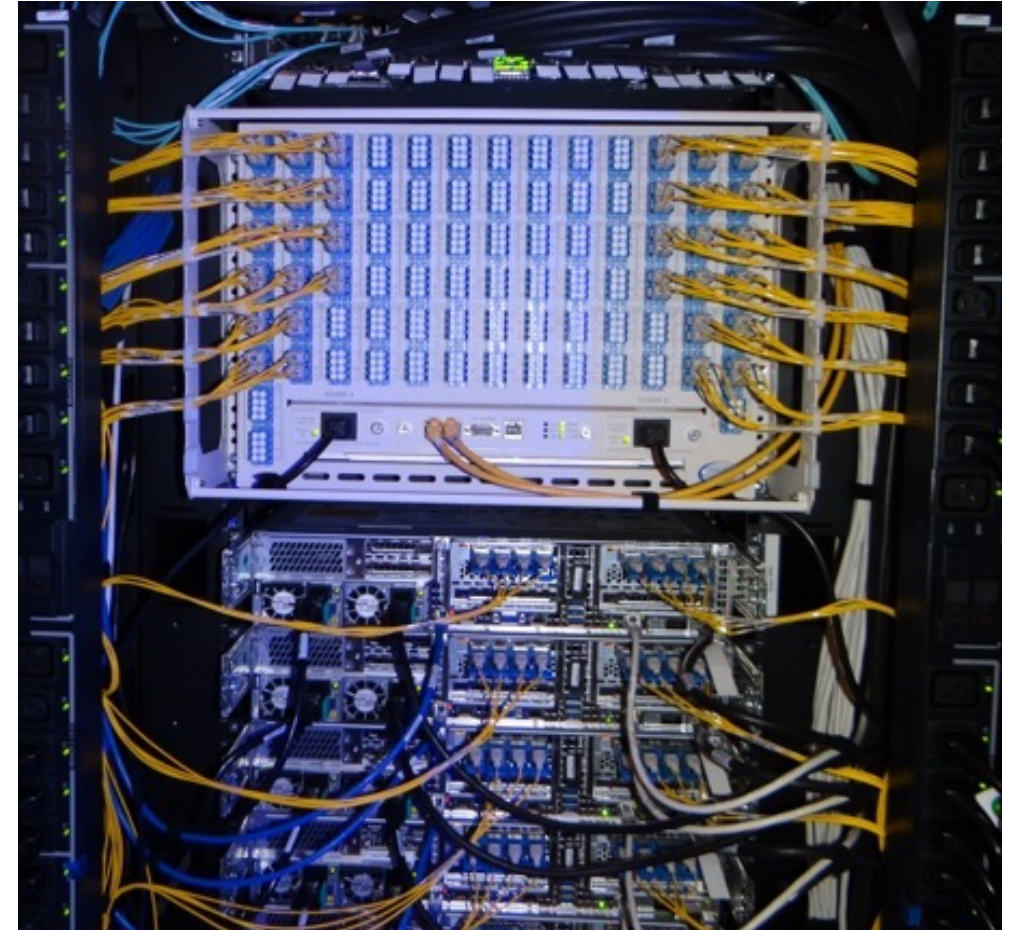
Supercomputer Noctua 2

- HPC applications need immense compute power
- Only achievable with high degree of parallelism
- Nowadays: Use of accelerators (GPUs / FPGAs) to increase efficiency
- **Noctua 2 @ PC2**
 - 1,124 servers
 - 143,872 CPU cores
 - 128 NVIDIA A100 GPUs
 - 100 Gb/s InfiniBand network
 - 48 Xilinx Alveo U280 FPGA accelerators
 - 32 BittWare 520N with Intel Stratix 10



Supercomputer Noctua 2

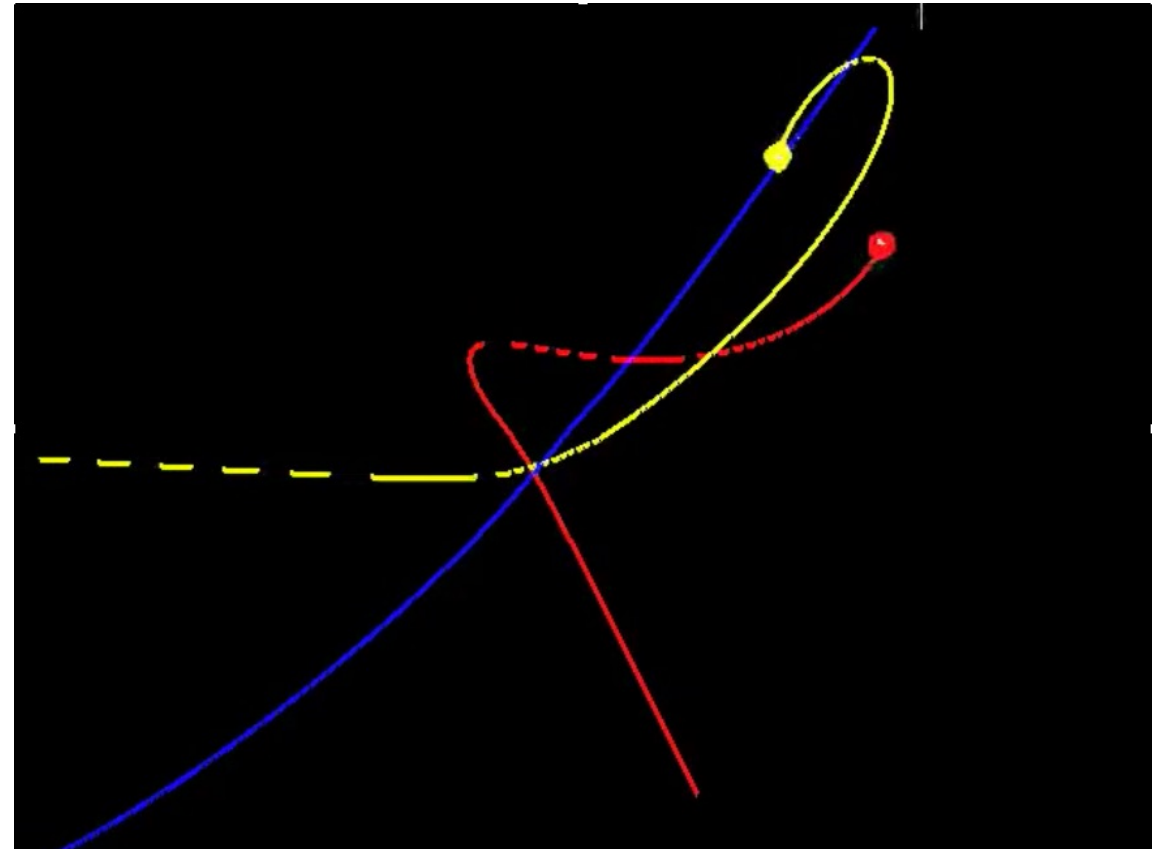
- Worldwide unique setup:
 - Top of the line FPGAs of both big vendors
 - 80 FPGAs within a single cluster
 - Direct optical network between FPGAs
 - Optical circuit switching for ultra low latency
 - Possibility for custom network topologies and protocols
- Chance to combine FPGA acceleration with parallelism
 - Multi-FPGA implementations



Successful Multi-FPGA Implementations

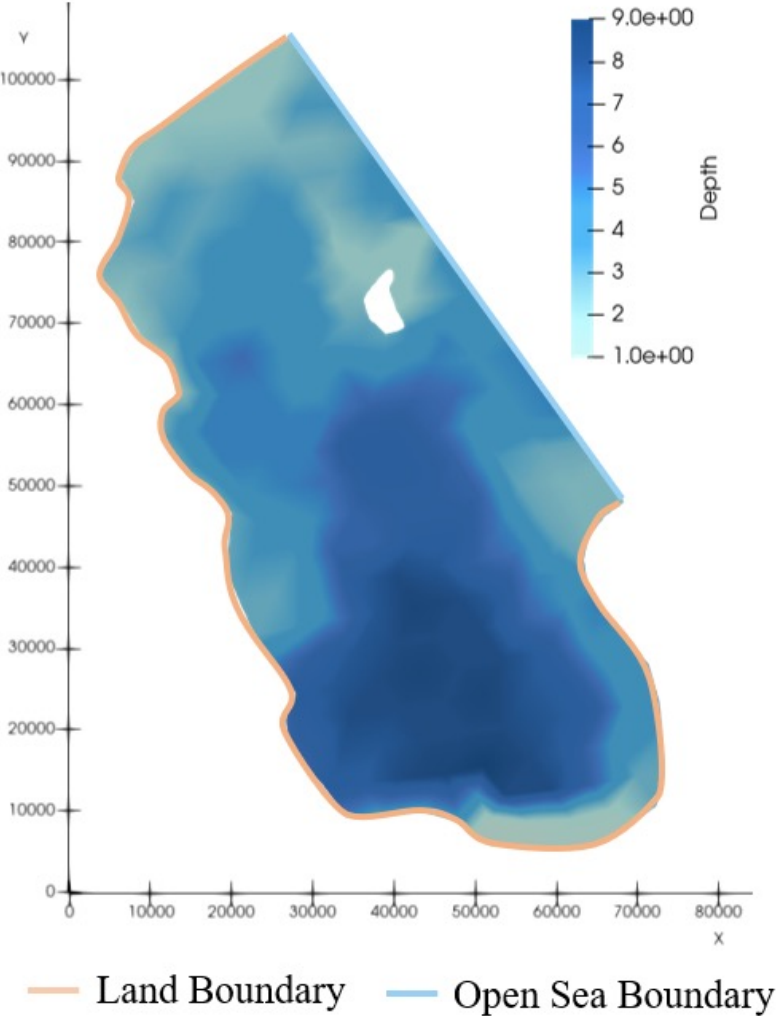
N-Body Simulation

- Simulation of moving objects
- Applicable to all kinds of scales:
 - from galaxy clusters in space
 - down to atoms
- Idea:
 - calculate forces acting between objects
 - update movement based on forces
 - repeat
- Implemented on 24 FPGAs
 - much better performance scaling than on CPUs

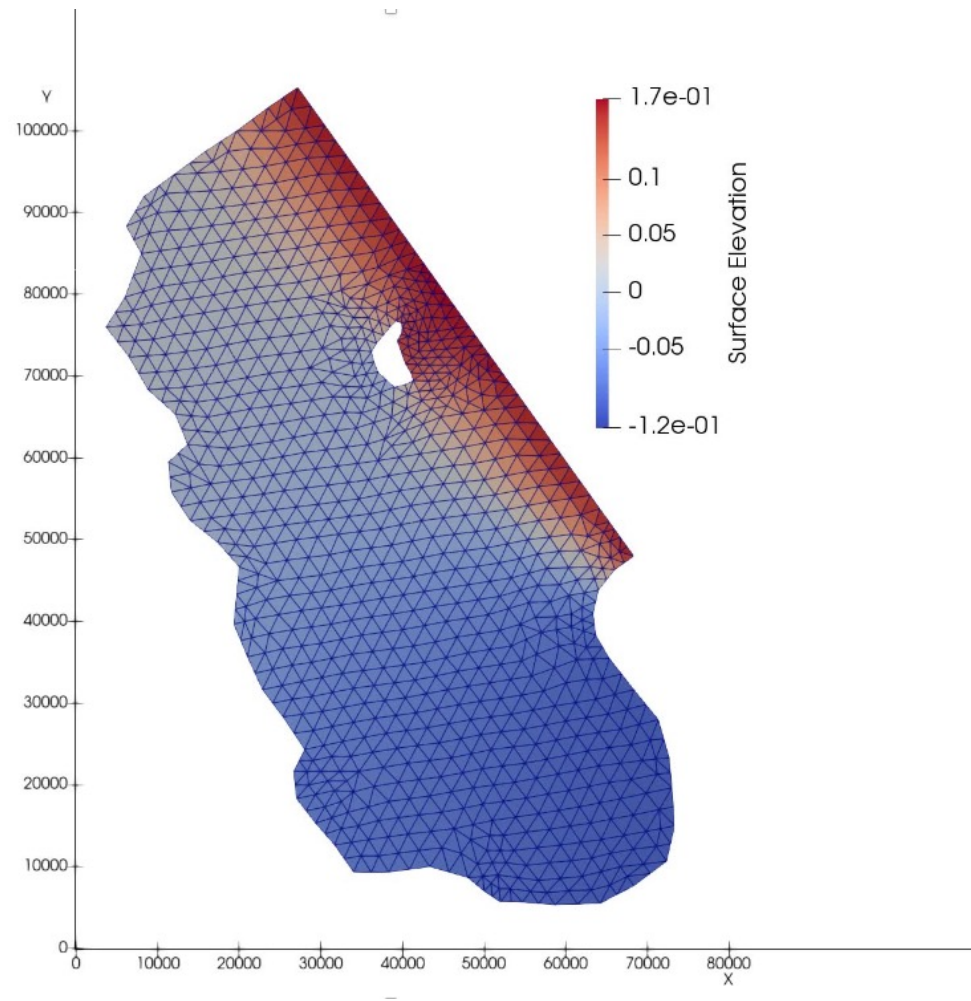


Shallow Water Simulations

- Simulation of water levels / waves near the coast



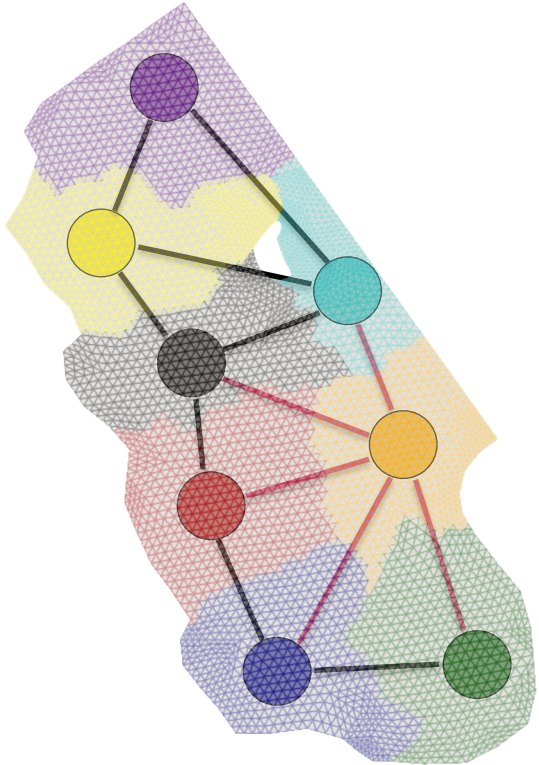
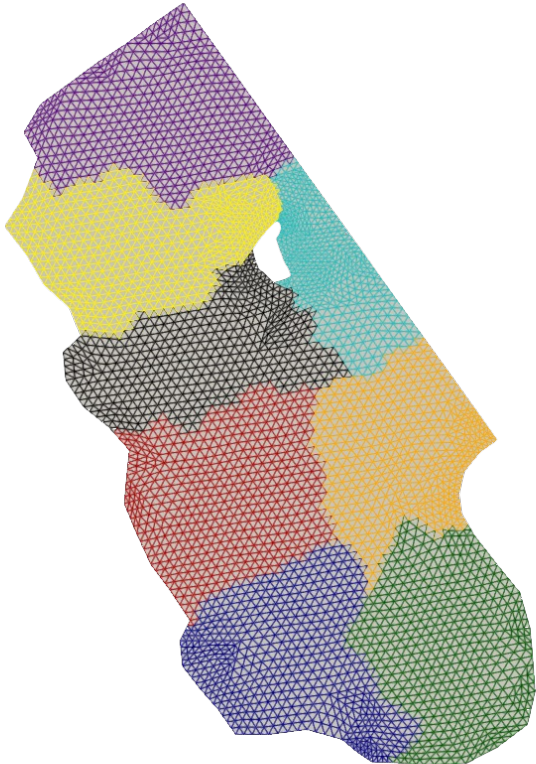
Shallow Water Simulations



- Simulation of water levels / waves near the coast
- Area divided into differently sized elements → “unstructured grid”

Shallow Water Simulations

- Grid partitions distributed among available FPGAs
- Near-perfect scaling with increasing number of FPGAs



Project Group MulFPGA

Goal of the Project Group

- Implement HPC methods or key algorithms on multiple networked FPGAs
- Examples for relevant application domains:
 - Numeric simulations on structured or unstructured grids
 - Dense or sparse linear algebra
 - Parallel graph processing



What you will do

- **Tutorial phase: Get familiar with FPGA programming**
- **Learn about performance modeling in HPC in general and for FPGA acceleration in particular**
- **Explore different communication modes for Multi-FPGA applications and characterize them**
- **Port one or two HPC applications to a Multi-FPGA design, testing different communication variants or using the analytically determined best one**

Requirements

- Interest in high-performance computing and accelerator technology
- Familiarity with computer architecture and networking, i.e., terms like
 - Pipelining
 - Parallelism
 - Delay, Bandwidth
- Some programming experience will help
 - FPGAs are programmed using Intel oneAPI (C++ extension), OpenCL, C, C++
 - *No* hardware description languages used (VHDL / Verilog)
 - You don't have to be an expert in C/C++ but should have the willingness to learn

Requirements

- Interest in high-performance computing and accelerator technology

- Familiarity

- Pipelining
- Parallelism
- Delay, Buffering

- Some projects

- FPGAs and ASICs
- No hardware description languages
- You don't need a PhD

```
q.submit([&](handler &h) {  
    auto acc_A = buffer_A.get_access<access::mode::read>(h);  
    auto acc_B = buffer_B.get_access<access::mode::read>(h);  
    auto acc_C = buffer_C.get_access<access::mode::write>(h);  
  
    h.single_task<class vector_sum>([=]() {  
  
        for (int i=0; i<10000; i++) {  
            acc_C[i] = acc_A[i] + acc_B[i];  
        }  
  
    });  
}).wait();
```

You will learn about...

- HPC applications and algorithms
- Modern HPC architectures
- Performance estimation / modelling
- Modern development flows for high-performance FPGAs
- Different networking approaches

Image removed for PDF export



Team & Contact



Prof. Christian Plessl

Tobias Kenter



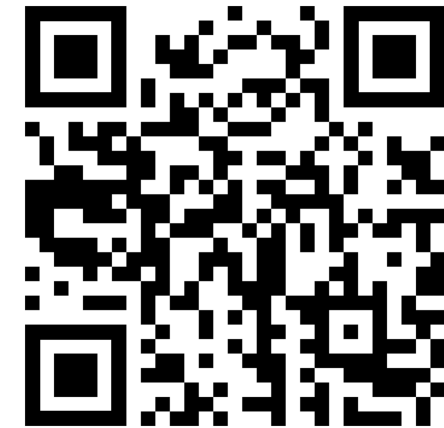
Michael Laß



Heinrich Riebler



Located in the X building



<https://en.cs.uni-paderborn.de/hpc/>